Better Streaming Algorithms for the Maximum Coverage Problem

Andrew McGregor*
University of Massachusetts
mcgregor@cs.umass.edu

Hoa T. Vu*
University of Massachusetts
hvu@cs.umass.edu

arXiv:1610.06199v2 [cs.DS] 21 Dec 2016

Abstract

We study the classic NP-Hard problem of finding the maximum k-set coverage in the data stream model: given a set system of m sets that are subsets of a universe $\{1, \ldots, n\}$, find the k sets that cover the most number of distinct elements. The problem can be approximated up to a factor $1 - 1/e$ in polynomial time. In the streaming-set model, the sets and their elements are revealed online. The main goal of our work is to design algorithms, with approximation guarantees as close as possible to $1 - 1/e$, that use sublinear space $o(mn)$. Our main results are:

- Two $(1 - 1/e - \epsilon)$ approximation algorithms: The first uses $O(\epsilon^{-1} \log(k/\epsilon))$ passes and $\tilde{O}(\epsilon^{-2}k)$ space[1] whereas the second uses only a single pass but $\tilde{O}(\epsilon^{-3}m)$ space.
- We show that any approximation factor better than $(1 - 1/e)$ in constant passes requires $\Omega(m)$ space for constant k even if the algorithm is allowed unbounded processing time. We also demonstrate a simple *single* pass and $(1 - \epsilon)$ approximation algorithm that uses $\tilde{O}(\epsilon^{-2}mk)$ space.

We also study the maximum k-vertex coverage problem in the dynamic graph stream model. In this model, the stream consists of edge insertions and deletions of a graph on N vertices. The goal is to find k vertices that cover the most number of distinct edges.

- We show that any constant approximation in constant passes requires $\Omega(N)$ space for constant k whereas $\tilde{O}(\epsilon^{-2}N)$ space is sufficient for a $(1 - \epsilon)$ approximation and arbitrary k in a single pass.
- For regular graphs, we show that $\tilde{O}(\epsilon^{-3}k)$ space is sufficient for a $(1 - \epsilon)$ approximation in a single pass. We generalize this to a $(\kappa - \epsilon)$ approximation when the ratio between the minimum and maximum degree is bounded below by κ.

*This work was supported by NSF Awards CCF-0953754, IIS-1251110, CCF-1320719, and a Google Research Award.
[1] $\tilde{O}(\cdot)$ suppresses polylog factors.

1 Introduction

The *maximum set coverage problem* is a classic NP-Hard problem that has a wide range of applications including facility and sensor allocation [32], information retrieval [6], influence maximization in marketing strategy design [28], and the blog monitoring problem where we want to choose a small number of blogs that cover a wide range of topics [39]. In this problem, we are given a set system of m sets that are subsets of a universe $[n] := \{1, \ldots, n\}$. The goal is to find the k sets whose union covers the largest number of distinct elements. For example, in the application considered by Saha and Getoor [39], the universe corresponds to n topics of interest to a reader, each subset corresponds to a blog that covers some of these topics, and the goal is to maximize the number of topics that the reader learns about if she can only choose k blogs. It is well-known that the greedy algorithm, which greedily picks the set that covers the most number of uncovered elements, is a $1 - 1/e$ approximation. Furthermore, unless $P = NP$, this approximation factor is the best possible [23].

The *maximum vertex coverage problem* is a special case of this problem in which the universe corresponds to the edges of a given graph and there is a set corresponding to each node of the graph that contains the subset of edges that are incident to that node. For this problem, algorithms based on linear programming are known to achieve a $3/4$ approximation for general graphs [1] and a $8/9$ approximation for bipartite graphs [14]. Assuming $P \neq NP$, there does not exist a polynomial-time approximation scheme. Recent work has focused on finding purely combinatorial algorithms for this problem [13].

Streaming Algorithms. Unfortunately, for both problems, the aforementioned greedy and linear programming algorithms do not scale well to massive data sets. This has motivated a significant research effort in designing algorithms that could handle large data in modern computation models such as the data stream model and the MapReduce model [10, 33]. In the data stream model, the k-set coverage problem and the related set cover problem have received a lot of attention in recent research [7, 9, 17, 22, 25, 43].

Two variants of the data stream model are relevant to our work. In the *streaming-set model* [22, 24, 30, 38, 39, 42], the stream consists of m sets $S_1, \ldots, S_m$ and each S_i is encoded as the list of elements in that set along with a unique ID for the set. In the dynamic graph stream model [2–5, 8, 12, 19, 24, 26, 27, 31, 35, 36], relevant to the maximum vertex coverage problem, the stream consists of insertions and deletions of edges of the underlying graph. For a recent survey of research in graph streaming, see [34]. Note that any algorithm for the dynamic graph stream model can also be used in the streaming-set model; the streaming-set model is simply a special case in which there is no deletion and edges are grouped by endpoint.

1.1 Related Work

Maximum Set Coverage. Saha and Getoor [39] gave a swap-based $1/4$ approximation algorithm that uses a single pass and $\tilde{O}(kn)$ space. At any point, their algorithm stores k sets explicitly in the memory as the current solution. When a new set arrives, based on a specific rule, their algorithm either swaps it with the set with the least contribution in the current solution or does nothing and moves on to the next set in the stream. Subsequently, Ausiello et al. [9] gave a slightly different swap-based algorithm that finds a $1/4 + O(1/k^2)$ approximation using one pass and the same space. Yu and Yuan [43] claimed an $\tilde{O}(n)$ space and single-pass algorithm with an approximation factor around 0.3 based on the aid of computer simulation.

Recently, Badanidiyuru et al. [10] gave a generic single-pass algorithm for maximizing a monotone submodular function on the stream's objects subject to the cardinality constraint that at most k objects are selected. Their algorithm guarantees a $1/2 - \epsilon$ approximation. At a high level, based on a rule that is different from [9, 39] and a guess of the optimal value, their algorithm decides if the next object (which is a set in our case) is added to the current solution. The algorithm stops when it reaches the end of the stream or when k objects have been added to the solution. In the k-set coverage problem, the rule requires knowing the coverage of the current solution. As a result, a careful adaptation to the k-set coverage problem uses $\tilde{O}(\epsilon^{-1}n)$ space. For constant ϵ, this result directly improves upon [9, 39]. Subsequently, Chekuri et al. [18] extended this work to non-monotone submodular function maximization under constraints beyond cardinality.

The set cover problem, which is closely related to the k-set coverage problem, has been studied in [7, 17, 22, 25, 39]. We refer to [7] for a comprehensive summary of results and discussion.

Maximum Vertex Coverage. The streaming k-vertex coverage problem was studied by Ausiello et al. [9]. They first observed that simply outputting the k vertices with highest degrees is a $1/2$ approximation; this can easily be done in the streaming-set model. The main results of their work were $\tilde{O}(kN)$-space algorithms that have better approximation for special types of graph. Their results include a 0.55 approximation for regular graphs and a 0.6075 approximation for regular bipartite graphs. Note that their paper only considered the streaming-set model whereas our results for maximum vertex coverage will consider the more challenging dynamic graph stream model.

1.2 Our Contributions

Maximum k-set coverage. Our main goal is to achieve the $1 - 1/e$ approximation that is possible in the non-streaming or offline setting.

- We present polynomial time data stream algorithms that achieve a $1 - 1/e - \epsilon$ approximation for arbitrarily small ϵ. The first algorithm uses $O(\epsilon^{-1} \log(k/\epsilon))$ passes and $\tilde{O}(\epsilon^{-2}k)$ space whereas the second algorithm uses one pass but $\tilde{O}(\epsilon^{-3}m)$ space. Given that k is likely to be relatively small in many applications, we consider both algorithms to be pass efficient but the first algorithm uses much less space at the cost of using more than one pass. We note that storing the solution itself requires $\Omega(k)$ space. Thus, we consider $\tilde{O}(\epsilon^{-2}k)$ space to be surprisingly space efficient.

- We show that $\Omega(m)$ space is required by any constant pass (randomized) algorithm to achieve an approximation factor better than $1 - 1/e$ with probability at least 0.99; this holds even if the algorithm is permitted exponential time. To the best of our knowledge, this is the first non-trivial space lower bound for this problem. However, with exponential time and $\tilde{O}(\epsilon^{-2}mk)$ space we observe that a $1 - \epsilon$ approximation is possible in a single pass.

For a slightly worse approximation, a $1/2 - \epsilon$ approximation in one pass can be achieved using $\tilde{O}(\epsilon^{-3}k)$ space. This follows by building on the result of Badanidiyuru et al. [10]. Finally, we design a $1/3 - \epsilon$ approximation algorithm for the budgeted maximum set coverage problem using one pass and $\tilde{O}(n)$ space. In this version, each set S has a cost $w(S)$ in the interval $[0, L]$. The goal is to find a collection of sets whose total cost does not exceed L that cover the most number of distinct elements. Khuller et al. [29] presented a polynomial time and $1 - 1/e$ approximation algorithm based on the greedy algorithm and an enumeration technique. Our results are summarized in Figure 1.

Shortly after our submission, in an independent work, Bateni et al. [11] also presented a single-pass, $\tilde{O}(\epsilon^{-3}m)$ space algorithm that finds a $1 - 1/e - \epsilon$ approximation for the maximum k-set coverage problem. We note that our approach also works in their *edge arrival* model in which the stream reveals the set-element relationships one at a time.

Maximum k-vertex coverage. Compared to the most relevant previous work [9], we study this problem in a more general model, i.e., the dynamic graph stream model. We manage to achieve a better approximation and space complexity for general graphs even when comparing to their results for special types of graph. Our results are summarized in Figure 2. In particular, we show that

- $\tilde{O}(\epsilon^{-2}N)$ space is sufficient for a $1 - \epsilon$ approximation (or a $3/4 - \epsilon$ approximation if restricted to polynomial time) and arbitrary k in a single pass. The algorithms in [9] use $\tilde{O}(kN)$ space and achieve an approximation worse than 0.61 even for special graphs.

- Any constant approximation in constant passes requires $\Omega(N)$ space for constant k.

- For regular graphs, we show that $\tilde{O}(\epsilon^{-3}k)$ space is sufficient for $1 - \epsilon$ approximation in a single pass. We generalize this to an $\kappa - \epsilon$ approximation when the ratio between the minimum and maximum degree is bounded below by κ. We also extend this result to hypergraphs.

Upper/Lower bound	Number of passes	Space	Approximation	Constraint
U	$O\left(\epsilon^{-1}\log(k/\epsilon)\right)$	$\tilde{O}(\epsilon^{-2}k)$	$1-1/e-\epsilon$	C
U	1	$\tilde{O}(\epsilon^{-3}k)$	$1/2-\epsilon$	C
U	1	$\tilde{O}(\epsilon^{-3}m)$	$1-1/e-\epsilon$	C
U	1	$\tilde{O}(\epsilon^{-2}mk)$	$1-\epsilon$	C
U	1	$\tilde{O}(\epsilon^{-1}n)$	$1/3-\epsilon$	B
L	constant	$\Omega(mk^{-2})$	$1-1/e+\epsilon$	C

Figure 1: Summary of results for MaxSetCoverage, C: cardinality, B: budgeted

Upper/Lower bound	Number of passes	Space	Approximation
U	1	$\tilde{O}(\epsilon^{-2}N)$	$1-\epsilon$
U	1	$\tilde{O}(\epsilon^{-3}k)$	$\kappa-\epsilon$
L	1	$\Omega(N)$	constant

Figure 2: Summary of results for MaxVertexCoverage
κ is the ratio between the minimum and maximum degrees

Our techniques. On the algorithmic side, our basic approach is a "guess, subsample, and verify" framework. At a high level, given any basic streaming algorithm, equipped with a near-correct guess of the optimal coverage, that finds an approximate optimal k-coverage in the original universe U, we could run the same algorithm on a subsampled universe U' defined by a carefully chosen hash function. The hash function depends on the guess value.

If the guess is at least nearly correct, running the algorithm on the subsampled universe results in a small space complexity. However, there are two main challenges. First, an algorithm instance with a wrong guess could use too much space. We simply terminate those instances. The second issue is more subtle. Because the hash function is not fully independent, we appeal to a special version Chernoff bound. The bound needs not guarantee a good approximation unless the guess is near-correct. To this end, we use the F_0 estimation algorithm to verify the coverage of the solutions. Finally, we return the solution with maximum estimate coverage. This framework allows us to restrict the analysis solely to the near-correct guess. The analysis is, therefore, significantly simpler.

Some of our other algorithmic ideas are inspired by previous works. The "thresholding greedy" technique was inspired by [17, 21]. However, the analysis is different for our problem. Furthermore, to optimize the number of passes, we rely on new observations.

Another algorithmic idea in designing one-pass space-efficient algorithm is to treat the sets differently based on their contributions. During the stream, we immediately add the sets with large contributions to the solution. In the contrary, we store the contribution of each remaining sets explicitly and solve the remaining problem offline. Har-Peled et al. [25] devised a somewhat similar strategy but the details are different.

For the k-vertex coverage problem, we show that simply running the streaming cut-sparsifier algorithm is sufficient and optimal up to a polylog factor. The novelty is to treat it as an interesting corner case of a more space-efficient algorithm for near regular graphs, i.e., κ is bounded below.

One of the novelties is proving the lower bound via a randomized reduction from the k-party set disjointness problem.

2 Algorithms for maximum k-set coverage

In this section, we design various algorithms for approximating MaxSetCoverage in the data stream model. Our main algorithmic results in this section are two $1-1/e-\epsilon$ approximation algorithms. The first uses $O(\epsilon^{-1}\log(k/\epsilon))$ passes but only $\tilde{O}(\epsilon^{-2}k)$ space whereas the second takes a single pass but uses more space. We also briefly explore some other trade-offs in a third section.

Preliminaries. An important component of many of the algorithms in this section is *universe-subsampling* where we essentially ignore many of the elements in each set and then run an approximation on the stream

of subsampled sets. To analyze this, we will need to appeal to the following Chernoff bounds for limited independent random variables.

Theorem 1 (Schmidt et al. [40]). *Let $X_1, \ldots, X_n$ be boolean random variables. Let $X = \sum_{i=1}^{n} X_i$ and $\mu = \mathrm{E}[X]$. Suppose $\mu \leq n/2$. If X_i are $\lceil \gamma\mu \rceil$-wise independent, then*

$$\Pr[|X - \mu| \geq \gamma\mu] \leq \exp\left(-\lfloor \min(\gamma, \gamma^2) \cdot \mu/3 \rfloor\right) .$$

Another important tool we rely on is F_0*-sketching*. For the purposes of our application, we can summarize the existing result as follows:

Theorem 2 (Cormode et al. [20]). *There exists an $\tilde{O}(\epsilon^{-2} \log \delta^{-1})$-space algorithm that, given a set $S \subseteq [n]$, can construct a data structure $\mathcal{M}(S)$, called an F_0* sketch *of S, that has the property that the number of distinct elements in a collection of sets $S_1, S_2, \ldots, S_k$ can be approximated up to a $1 \pm \epsilon$ factor with probability at least $1 - \delta$ given the collection of F_0 sketches $\mathcal{M}(S_1), \mathcal{M}(S_2), \ldots, \mathcal{M}(S_k)$.*

Throughout this section, we will frequently use some simple inequalities including that for all $k \geq 1$, we have $1-(1-1/k)^k \geq 1-1/e$ and for all $\epsilon \geq 0$, we have $\log_{1+\epsilon} n = O(\epsilon^{-1} \log n)$ and $1 - e^{-1/(1+\epsilon)} \geq 1 - 1/e - \epsilon$.

Notations. For convenience, we write $z \approx_\epsilon x$ if $(1-\epsilon)x \leq z \leq (1+\epsilon)x$. Additionally, we write $z \approx_{\epsilon^+} x$ if $x \leq z \leq (1+\epsilon)x$ and finally $z \approx_{\epsilon^-} x$ if $(1-\epsilon)x \leq z \leq x$. If $\mathcal{A}$ is a collection of sets, then $\mathcal{C}(\mathcal{A})$ denotes the union of these sets.

2.1 $(1 - 1/e - \epsilon)$ approximation in $O(\epsilon^{-1} \log(k\epsilon^{-1}))$ passes and $\tilde{O}(\epsilon^{-2}k)$ space

Approach. The algorithm is based on the standard greedy approach but instead of adding the set that increases the coverage of the current solution the most at each set, we add a set if the number of new elements covered by this set exceeds a certain threshold. This threshold decreases with each pass in such a way that after only $O(\epsilon^{-1} \log(k/\epsilon))$ passes, we have a good approximate solution but the resulting algorithm may use too much space. We will fix this by first randomly subsampling each set at different rates and running multiple instantiations of the basic algorithm corresponding to different rates of subsampling.

Technical Details. For the analysis, it will be convenient to use $z = \mathrm{OPT}$ to denote the maximum number of elements that can be covered by k sets. We start by designing a $(1 - 1/e - \epsilon)$ approximation algorithm that uses $\tilde{O}(k + z)$ space and $O(\epsilon^{-1} \log(k/\epsilon))$ passes. See Algorithm 1. We will subsequently use universe-subsampling to reduce the space to $\tilde{O}(\epsilon^{-2}k)$.

Lemma 3. *There exists an $O(\epsilon^{-1} \log(k/\epsilon))$-pass, $O(k \log m + z \log n)$-space algorithm (Algorithm 1) that finds a $1 - 1/e - \epsilon$ approximation of* MaxSetCoverage.

To analyze the algorithm we introduce some notation. After the ith set was picked, let a_i be the number of new elements covered by this set and let b_i be the total number of covered elements so far. Furthermore, let $c_i = z - b_i$. We define $a_0 := 0$, $b_0 := 0$, and $c_0 := z$.

Lemma 4. *For $0 \leq i \leq k-1$, $a_{i+1} \geq c_i/(\alpha k)$.*

Proof. Suppose the algorithm picked the $(i+1)$th set S during the jth pass. Let C be the set of covered elements just before the algorithm added S to the solution Sol.

Note that just before the algorithm added S, there must exist a set S' (which could be S) that was not in Sol and $|S' \setminus C| \geq c_i/k$. This follows because the optimum collection of k sets covers at least c_i elements that are currently uncovered and hence one of these sets must cover at least c_i/k new elements. On the other hand, since S' was not already in Sol, we know that S' was not added during the first $j-1$ passes. Hence, $|S' \setminus C| < \tau/\alpha^{j-1}$; note that if $j = 1$, $|S' \setminus C| \leq \tau/\alpha^{1-1}$ is trivially satisfied because $|S'| \leq \tau$. Therefore,

$$\tau/\alpha^{j-1} \geq |S' \setminus C| \geq c_i/k$$

and in particular, $\tau/\alpha^j \geq c_i/(\alpha k)$.

Since the algorithm picked S, we have $a_{i+1} = |S \setminus C| \geq \tau/\alpha^j$ and the lemma follows because we have just concluded that $\tau/\alpha^j \geq c_i/(\alpha k)$. □

$C \leftarrow \emptyset$.
Let $\alpha = 1 + \epsilon$, $\beta = \lceil \log_\alpha(k/\epsilon) \rceil$.
Let τ be the size of the largest set (this can be found in one extra pass if necessary).
for $j \leftarrow 0$ **to** β **do**
  Take a pass over the stream. On seeing a set S:
  if $|S \setminus C| \geq \tau/\alpha^j$ **then**
    $\text{Sol} \leftarrow \text{Sol} \cup \{S\}$. `// Sol is only used in the analysis`
    $\text{ID}_{\text{Sol}} \leftarrow \text{ID}_{\text{Sol}} \cup \{\text{ID}(S)\}$.
    $C \leftarrow C \cup S$.
    Store the F_0 sketch of S.
    if $|\,\text{Sol}\,| = k$ **then return** ID_{Sol} and stop the algorithm.
  end
end
Post-processing: **return** ID_{Sol}.

Algorithm 1: The basic algorithm for a $(1 - 1/e - \epsilon)$ approximation in $O(\epsilon^{-1} \log(k/\epsilon))$ passes and $\tilde{O}(\epsilon^{-2}k)$ space.

Proof of Lemma 3. To argue about the approximation factor, we first prove that we always have $c_i \leq \left(1 - \frac{1}{\alpha k}\right)^i z$. Obviously, we have $c_0 \leq (1 - \frac{1}{\alpha k})^0 z$. Suppose $c_i \leq (1 - \frac{1}{\alpha k})^i z$. Then, according to Lemma 4, $a_{i+1} \geq \frac{c_i}{\alpha k}$. Thus,

$$c_{i+1} = c_i - a_{i+1} \leq c_i - \frac{c_i}{\alpha k} = c_i \left(1 - \frac{1}{\alpha k}\right) \leq z \left(1 - \frac{1}{\alpha k}\right)^{i+1}.$$

To simplify the rest of the argument, suppose we had set $\beta = \lceil \log_\alpha \tau \rceil$. After the last iteration where $j = \beta$, we either have a solution of k sets or we already cover the entire universe with fewer than k sets. We only need to consider the first case in which

$$c_k \leq \left(1 - \frac{1}{\alpha k}\right)^k z \leq e^{-1/\alpha} z \leq (1/e + \epsilon) z.$$

Thus, if we run the algorithm with $\beta = \lceil \log_\alpha \tau \rceil$, the solution's coverage is at least

$$b_k \geq z - c_k \geq (1 - 1/e - \epsilon) z.$$

We call this hypothetical solution $\widetilde{\text{Sol}}$. We observe that in this case the number of passes could be $\lceil \log_\alpha \tau \rceil$ which is $\Theta(\epsilon^{-1} \log n)$ in the worst case. To optimize the number of passes, we continue the argument as follows.

Because we actually set $\beta = \lceil \log_\alpha(k/\epsilon) \rceil$, any set that is in $\widetilde{\text{Sol}}$ and not in the actual solution must be added after the pass where $j = \lceil \log_\alpha(k/\epsilon) \rceil$; therefore, it must cover fewer than $\frac{\tau}{k/\epsilon}$ new elements. Thus, the number of elements in the solution after the $(\lceil \log_\alpha(k/\epsilon) \rceil)$th pass is at least

$$|\,\mathcal{C}(\widetilde{\text{Sol}})| - \epsilon\tau \geq (1 - 1/e - \epsilon)\,\text{OPT} - \epsilon\,\text{OPT} \geq (1 - 1/e - 2\epsilon)\,\text{OPT}\,.$$

Regarding the space and passes used by the algorithm: It is immediate that the number of passes is $O(\log_\alpha(k/\epsilon))$. The algorithm only needs to store the IDs and F_0 sketches of the sets in the solution and the set C of covered elements so far. Since $|C| \leq z$, we conclude the algorithm uses $O(k \log m + z \log n)$ space. □

Reducing the space via element subsampling. Next, we show that an appropriate subsampling approach improves the space complexity. Let v be a guess for OPT that satisfies $\text{OPT}/2 \leq v \leq \text{OPT}$.

Let c be some sufficiently large constant. We assume that $k \leq \epsilon^2 n/(8c \log m)$. Otherwise, appealing to Lemma 3, we deduce an algorithm that uses $O(n \log n + k \log m) = O(\epsilon^{-2} k \log m \log n + k \log m) = \tilde{O}(\epsilon^{-2}k)$ space and $O(\epsilon^{-1} \log(k/\epsilon))$ passes as required.

Let $h : [n] \to \{0,1\}$ be drawn from a family of $(2c\epsilon^{-2}k\log m)$-wise independent hash functions where

$$p := \Pr[h(e) = 1] = \frac{c\epsilon^{-2}k\log m}{v}.$$

The space to store h is $\tilde{O}(\epsilon^{-2}k)$. For any set S that is a subset of $[n]$, we define

$$S' := \{e \in S : h(e) = 1\} .$$

The next lemma and its corollary will allow us to argue that approximating the maximum coverage amongst the elements $\{e \in [n] : h(e) = 1\}$ gives only a slightly weaker approximation of the maximum coverage amongst the original set of elements.

Lemma 5. *Assuming $k \le \epsilon^2 n/(8c\log m)$, with high probability[2], for all collections of k sets $S_1, \ldots, S_k$ in the stream:*

$$|S'_1 \cup \ldots \cup S'_k| = |S_1 \cup \ldots \cup S_k|p \pm \epsilon vp .$$

Proof. Fix any collection of k sets $S_1, \ldots, S_k$. Let $D = |S_1 \cup \ldots \cup S_k|$ and $D' = |S'_1 \cup \ldots \cup S'_k|$. We first observe that

$$\mu := \mathrm{E}[D'] = pD \le p\,\mathrm{OPT} < 2pv = 2c\epsilon^{-2}k\log m \le n/2.$$

Appealing to the Chernoff bound with limited independence, Theorem 1, we have

$$\Pr[|D' - \mu| \ge \epsilon vp] = \Pr[|D' - \mu| \ge \gamma Dp] \le \exp\left(-\lfloor \min(\gamma, \gamma^2)\cdot \mu/3\rfloor\right)$$

where $\gamma = \epsilon v/D$ since the hash function was $\lceil\gamma\mu\rceil = \lceil\epsilon vp\rceil$-wise independent. But note that

$$\begin{aligned}\exp\left(-\lfloor \min(\gamma, \gamma^2)\cdot \mu/3\rfloor\right) &= \exp\left(-\lfloor \min(1, \gamma)\cdot \epsilon vp/3\rfloor\right)\\ &\le \exp\left(-\lfloor 1/2\cdot ck\log m/3\rfloor\right) \le 1/m^{10k}\end{aligned}$$

where we use the fact that $\gamma = \epsilon v/D \ge \epsilon/2$ because $D \le \mathrm{OPT} \le 2v$. The lemma follows by taking the union bound over all $\binom{m}{k}$ collections of k sets. □

In particular, the following corollary establishes that a $1/t$ approximation of MaxSetCoverage when restricted to elements in $\{e \in [n] : h(e) = 1\}$ yields a $(1/t - 2\epsilon)$ approximation and at most $p\,\mathrm{OPT}(1+\epsilon) = O(\epsilon^{-2}k\log m)$ of these elements can be covered by k sets.

Corollary 6. *Assuming $k \le \epsilon^2 n/(8c\log m)$ and let OPT' be optimum number of elements that can be covered from $\{e \in [n] : h(e) = 1\}$. Then,*

$$p\,\mathrm{OPT}(1+\epsilon) \ge \mathrm{OPT}' \ge p\,\mathrm{OPT}(1-\epsilon)$$

Furthermore if $U_1, \ldots, U_k$ satisfies $|U'_1 \cup \ldots \cup U'_k| \ge p\,\mathrm{OPT}(1-\epsilon)/t$ for $t \ge 1$ then

$$|U_1 \cup \ldots \cup U_k| \ge \mathrm{OPT}(1/t - 2\epsilon) .$$

Proof. The fact that $\mathrm{OPT}' \ge p\,\mathrm{OPT}(1-\epsilon)$ follows by applying Lemma 5 to the optimum solution. According to Lemma 5, for all collections of k sets $U_1, \ldots, U_k$, we have $|U'_1 \cup \ldots \cup U'_k| = |U_1 \cup \ldots \cup U_k|p \pm \epsilon vp \le p\,\mathrm{OPT}(1+\epsilon)$ which implies the first inequality.

Now, suppose $|U'_1 \cup \ldots \cup U'_k| \ge p\,\mathrm{OPT}(1-\epsilon)/t$. Since $|U'_1 \cup \ldots \cup U'_k| - \epsilon vp \le |U_1 \cup \ldots \cup U_k|p$, we deduce that

$$|U_1 \cup \ldots \cup U_k| \ge \mathrm{OPT}(1-\epsilon)/t - \epsilon v \ge \mathrm{OPT}(1/t - 2\epsilon).$$

□

[2]We consider $1 - 1/\operatorname{poly}(m)$ or $1 - 1/\operatorname{poly}(n)$ as high probability.

Hence, if we know v such that $\mathrm{OPT}/2 \leq v \leq \mathrm{OPT}$, then we can run Algorithm 1 on the subsampled stream; the algorithm returns a $1 - 1/e - \epsilon$ on the stream of subsampled sets and hence a $1 - 1/e - 2\epsilon$ approximation to the original problem. Furthermore, the space used by the algorithm is $O(\epsilon^{-2}k \log m \log n)$ since for the subsampled sets, the maximum number of elements that can be covered by k sets is $O(\epsilon^{-2}k \log m)$.

To get around the fact that we do not know v, we simply run $\log n$ copies of the process with guesses $v = 1, 2, 4, 8, \ldots n$ in parallel. If any of these copies ever uses more than $O(\epsilon^{-2}k \log m \log n)$ space we can terminate that copy since the corresponding guess must be too small. At least one copy will not terminate with high probability, the copy corresponding to the good guess. Amongst the copies that do not terminate, we do not need to determine which copy corresponded to the guess that was within a factor 2 of OPT. Instead, since each computes the F_0 sketches of the sets selected by the algorithm, we can simple use the sets collected by the algorithm that covers the most elements. By re-parameterizing $\epsilon \leftarrow \epsilon/3$, we have therefore established the following theorem.

Theorem 7. *There exists an algorithm using $\tilde{O}(\epsilon^{-2}k)$ space and $O(\epsilon^{-1}\log(k\epsilon^{-1}))$ passes that finds a $1 - 1/e - \epsilon$ approximation of* MaxSetCoverage *with high probability.*

2.2 $(1 - 1/e - \epsilon)$ approximation in one pass and $\tilde{O}(\epsilon^{-3}m)$ space

The algorithm is provided $z \approx_{\epsilon^-} \mathrm{OPT}$.
$r \leftarrow z$, $C \leftarrow \emptyset$.
During the stream, on seeing a set S:
  Store F_0 sketch of S.
  if $|S \setminus C| \geq r/k$ **then**
    $\mathrm{Sol} \leftarrow \mathrm{Sol} \cup \{S\}$. `// Sol is only used for the analysis`
    $\mathrm{ID}_{\mathrm{Sol}} \leftarrow \mathrm{ID}_{\mathrm{Sol}} \cup \{\mathrm{ID}(S)\}$.
    $C \leftarrow C \cup S$.
    $r \leftarrow r - |S \setminus C|$.
    if $|\mathrm{Sol}| = k$ **then return** $\mathrm{ID}_{\mathrm{Sol}}$ and stop the algorithm.
  else
    $\mathcal{W} \leftarrow \mathcal{W} \cup \{\tilde{S}\}$ where $\tilde{S} = S \setminus C$.
  end
Post-processing:
  Use the greedy algorithm to find the remaining sets from $\mathcal{W}$. Store their IDs in $\mathrm{ID}_{\mathrm{Sol}2}$.
  $\mathrm{ID}_{\mathrm{Sol}} \leftarrow \mathrm{ID}_{\mathrm{Sol}} \cup \mathrm{ID}_{\mathrm{Sol}2}$.
  return $\mathrm{ID}_{\mathrm{Sol}}$.

Algorithm 2: The basic algorithm for a $1 - 1/e - \epsilon$ approximation in one pass and $\tilde{O}(\epsilon^{-3}m)$ space.

Approach. Our second algorithm also adds sets to the current solution if the number of new elements they cover exceeds some threshold. However, in contrast to the previous algorithm, this threshold does not remain constant throughout a pass and the algorithm only requires one pass. The basic algorithm relies on an estimate z of the optimum coverage OPT. Using this, the algorithm maintains a value r corresponding to a bound on the number of additional elements that could be covered by sets that are not in the current solution. The threshold for including a new set in the solution is that it covers at least r/k new elements. Unfortunately, this threshold is too high to ensure that we selected sets that achieve the required $1 - 1/e - \epsilon$ approximation and we may want to revisit adding a set, say S, that was not added when it first arrived. To facilitate this, we will explicitly store a set $\tilde{S}$ in $\mathcal{W}$ that includes only the elements of S that were uncovered when S arrived. By the fact that S was not added immediately, we know $\tilde{S}$ is not too large. At the end of the pass, we continue augmenting out current solutions using the collection $\mathcal{W}$. See Algorithm 2.

Technical Details. For the time being, we suppose that the algorithm is provided with an estimate $z \approx_{\epsilon^-} \mathrm{OPT}$. The algorithm uses C to keep track of the elements that have been covered so far. Upon seeing a

new set S, the algorithm stores $S \setminus C$ explicitly in $\mathcal{W}$ if S covers few new elements. Otherwise, the algorithm adds S to the solution and updates C immediately. At the end of the stream, if there are fewer than k sets in the solution, we use the greedy approach to find the remaining sets from $\mathcal{W}$.

Lemma 8. *Given an estimate $z \approx_{\epsilon^-} \mathrm{OPT}$, there exists a single-pass, $O\left(\frac{mz \log n}{k}\right)$-space algorithm (Algorithm 2) that finds a $1 - 1/e - \epsilon$ approximation of* MaxSetCoverage.

Proof. We observe that storing the set of covered elements C requires at most $\mathrm{OPT} \log n = O(z \log n)$ bits of space. For each set S such that $\tilde{S}$ is stored explicitly in $\mathcal{W}$, we need $O\left(\frac{z \log n}{k}\right)$ bits of space. Thus, the total space is $O\left(\frac{mz \log n}{k}\right)$.

After the algorithm added the ith set S to the solution, let a_i be the number of new elements that S covers and b_i be the total number of covered elements so far. Furthermore, for $i > 0$, let $c_i = z - b_i$. Define $a_0 := b_0 := 0$ and $c_0 := z$.

We can prove by induction that, $c_i \leq (1 - 1/k)^i z$. The base case where $i = 0$ is clearly true. Consider a set S that was added to Sol during the stream. Suppose S is the $(i + 1)$th set that was added to Sol. By the induction hypothesis, $c_i \leq (1 - 1/k)^i z$. Suppose $r = r'$ right before the algorithm added S. Since $a_{i+1} \geq r'/k = c_i/k$, we have

$$c_{i+1} = c_i - a_{i+1} \leq c_i - c_i/k \leq z(1 - 1/k)^{i+1} .$$

At the end of the stream, suppose $|\,\mathrm{Sol}\,| = j$. Then, $c_j \leq z(1 - 1/k)^j$. Consider the sets that were added to Sol offline after the stream. For $i \geq j$, we proceed with the usual greedy argument. Before the algorithm added the $(i + 1)$th set, there must be a set that covers at least c_i/k new elements. This is true because $z \approx_{\epsilon^-} \mathrm{OPT}$. Therefore, $c_{i+1} = c_i - a_{i+1} \leq c_i(1 - 1/k) \leq z(1 - 1/k)^{i+1}$. It is then straightforward to deduce the approximation by considering $i = k$.

□

If $k > \epsilon^2 n/(8c \log m)$ then according to Lemma 8, given a guess $z \approx_{\epsilon^-} \mathrm{OPT}$, we have a single-pass algorithm that uses $O(mz \log(n)/k) = \tilde{O}(\epsilon^{-2} m)$ space since $z \leq n < 8c\epsilon^{-2} k \log m$. We run $\lceil \log_{1+\epsilon} n \rceil = O(\epsilon^{-1} \log n)$ different copies of the algorithm corresponding to the guesses $z = (1 + \epsilon)^0, (1 + \epsilon)^1, (1 + \epsilon)^2, \ldots$ One of those guesses must be between $(1 - 2\epsilon)\,\mathrm{OPT}$ and OPT and hence the corresponding copy would return a $1 - 1/e - 2\epsilon$ approximation.

On the other hand, if $k \leq \epsilon^2 n/(8c \log m)$, we can then reuse the subsampling strategy in the previous section with one small change. Previously, it was sufficient to have a guess v that is within a factor of 2 of OPT. Now, since the algorithm requires knowledge of OPT, we need to use $O(\epsilon^{-1} \log n)$ guesses $v = (1 + \epsilon)^0, (1 + \epsilon)^1, (1 + \epsilon)^2, \ldots$ Therefore, there must be a guess $v \approx_{2\epsilon} \mathrm{OPT}$. According to Corollary 6,

$$\mathrm{OPT}' \approx_\epsilon p\,\mathrm{OPT} = (1 \pm \epsilon)\frac{c\epsilon^{-2} k \log m}{(1 \pm 2\epsilon)\,\mathrm{OPT}}\,\mathrm{OPT} \approx_{4\epsilon} c\epsilon^{-2} k \log m.$$

As a result, by setting $z = (1 - \epsilon)(1 - 4\epsilon)c\epsilon^{-2} k \log m$, we ensure that $z \approx_{(5\epsilon)^-} \mathrm{OPT}'$ when $v \approx_{2\epsilon} \mathrm{OPT}$. Applying Lemma 8 and Corollary 6 to the subsampled elements corresponding to the guess $v \approx_{2\epsilon} \mathrm{OPT}$, we have an algorithm copy that uses $\tilde{O}(\epsilon^{-2} m)$ space and finds a $1 - 1/e - O(\epsilon)$ approximation. Re-parameterizing ϵ results in the following theorem.

Theorem 9. *There exists a single-pass, $\tilde{O}(\epsilon^{-3} m)$ space algorithm that finds a $1 - 1/e - \epsilon$ approximation of* MaxSetCoverage *with high probability.*

2.3 Other Algorithmic Results

In this final subsection, we briefly review some other algorithmic results for MaxSetCoverage, either with different trade-offs or for a "budgeted" version of the problem.

$(1-\epsilon)$ **approximation in one pass and** $\tilde{O}(\epsilon^{-2}mk)$ **space.** In the previous subsection, we gave a single-pass $1-1/e-\epsilon$ approximation using $\tilde{O}(\epsilon^{-3}m)$ space. Here we observe that if we are permitted $\tilde{O}(\epsilon^{-2}mk)$ space and unlimited post-processing time then a $1-\epsilon$ approximation can be achieved directly from the F_0 sketches.

Specifically, in one pass we construct the F_0 sketches of all m sets, $\mathcal{M}(S_1), \ldots, \mathcal{M}(S_m)$ where the failure probability of the sketches is set to $\delta = 1/(nm^k)$. Thus, at the end of the stream, one can $1+\epsilon$ approximate the coverage $|S_{i_1} \cup \ldots \cup S_{i_k}|$ for each collection of k sets $S_{i_1}, \ldots, S_{i_k}$ with probability at least $1-1/(nm^k)$. Since there are at most $\binom{m}{k} \leq m^k$ collections of k sets, appealing to the union bound, we could guarantee that the coverages of all of the collections of k sets are preserved up to a $1+\epsilon$ factor with probability at least $1-1/n$. The space to store the sketches is $\tilde{O}(\epsilon^{-2}mk)$.

Theorem 10. *There exists a single-pass,* $\tilde{O}(\epsilon^{-2}mk)$*-space algorithm that finds a* $1-\epsilon$ *approximation of* MaxSetCoverage *with high probability .*

$(1/2-\epsilon)$ **approximation in one pass and** $\tilde{O}(\epsilon^{-3}k)$ **space.** We next observe that it is possible to achieve a $1/2-\epsilon$ using a single pass and $\tilde{O}(\epsilon^{-3}k)$ space. This follows immediately by combining the subsampling approach with a $1/2-\epsilon$ approximation of Badanidiyuru et al. [10] for the more general problem of streaming submodular maximization.

Theorem 11. *There exists a single-pass,* $\tilde{O}(\epsilon^{-3}k)$ *space algorithm that finds a* $1/2-\epsilon$ *approximation of* MaxSetCoverage *with high probability.*

Further details can be found in the appendix.

Budgeted Maximum Coverage. In this variation, each set S has a cost $w(\{S\}) \in [0, L]$. The problem asks to find the collection of sets whose total cost is at most L that covers the most number of distinct elements.

Given an estimate $z \approx_{\epsilon^-}$ OPT, we design a single-pass, $\tilde{O}(n)$-space and $1/3-\epsilon$ approximation algorithm. Our ideas and analysis are inspired by [10]. However, the algorithm, as well as the analysis, has to handle new corner cases and requires extra observations.

We consider Algorithm 3. We assume the cost of a set is given upon seeing it. Upon seeing a set S, based on a specific rule, S is either selected or discarded. If S is selected and adding S to Sol does not violate the budget constraint, the algorithm adds S to Sol. Otherwise, the algorithm outputs $\{S\}$ or the current solution Sol as the solution depending on which has the larger coverage.

The algorithm is provided $z \approx_{\epsilon^-}$ OPT.
$C \leftarrow \emptyset$, Sol $\leftarrow \emptyset$.
During the stream, on seeing a set S:
  if $|S \setminus C|/w(\{S\} \geq (2z/3 - |C|)/(L - w(\text{Sol}))$ **then**
    if $w(\text{Sol}) + w(\{S\}) \leq L$ **then**
      Sol $\leftarrow$ Sol$\cup\{S\}$. `// Sol is only used for the analysis`
      $\text{ID}_{\text{Sol}} \leftarrow \text{ID}_{\text{Sol}} \cup \{\text{ID}(S)\}$.
      $C \leftarrow C \cup S$.
    else
      return $\begin{cases} \text{ID}_{\text{Sol}} & \text{if } |C| \geq |S| \\ \{\text{ID}(S)\} & \text{otherwise} \end{cases}$ and then stop the algorithm. Call this step $(*)$.
    end
  end
Post-processing: **return** ID_{Sol}.

Algorithm 3: A single-pass, $\tilde{O}(n)$-space and $1/3-\epsilon$ approximation algorithm for the budgeted version

Lemma 12. *We always have* $|\mathcal{C}(\text{Sol})| \geq 2z \cdot w(\text{Sol})/(3L)$.

Proof. Let U_i be the ith set being added to Sol. We prove by induction on i. The base case $i = 1$ can be verified easily. Consider the set U_{i+1}. For simplicity, let $S = U_{i+1}$. From the induction hypothesis, just before S is added to Sol, we have $|C| \geq 2z \cdot w(\text{Sol})/(3L)$. We need to show that $|C \cup S| \geq 2z(w(\text{Sol}) + w(S))/(3L)$. Since the algorithm selected S, we have

$$\frac{|S \setminus C|}{w(\{S\})} \geq \frac{2z/3 - |C|}{L - w(\text{Sol})} .$$

Substituting $|S \setminus C| = |S \cup C| - |C|$, this implies

$$\begin{aligned}
|S \cup C| &\geq |C| + \frac{2z/3 - |C|}{L - w(\text{Sol})} \cdot w(S) \\
&= |C| \left(1 - \frac{w(S)}{L - w(\text{Sol})}\right) + \frac{2z/3 \cdot w(S)}{L - w(\text{Sol})} \\
&\geq \frac{2z \cdot w(\text{Sol})}{3L} \left(1 - \frac{w(S)}{L - w(\text{Sol})}\right) + \frac{2z/3 \cdot w(S)}{L - w(\text{Sol})} \\
&= \frac{2z(w(\text{Sol}) + w(S))}{3L},
\end{aligned}$$

where the last inequality follows from the induction hypothesis. □

Lemma 13. *If step $(*)$ in Algorithm 3 never occurred, then the final solution Sol_f satisfies*

$$|\,\mathcal{C}(\text{Sol}_f)| \geq z/3.$$

Proof. Suppose step $(*)$ never occurred. Consider a set S_j in the optimum solution A that was not selected. Consider C and Sol just before S_j was not selected. According to the algorithm's logic and Lemma 12, the set $\tilde{S}_j = S_j \setminus \mathcal{C}(\text{Sol}_f)$ of elements privately covered by S_j relative to $\mathcal{C}(\text{Sol}_f)$ satisfies:

$$\left|\tilde{S}_j\right| \leq |S_j \setminus C| \leq \left(\frac{2z}{3} - \frac{2z \cdot w(\text{Sol})}{3L}\right) \cdot \frac{w(S_j)}{L - w(\text{Sol})} = \frac{2z \cdot w(S_j)}{3L} .$$

Hence, we deduce that

$$\sum_{S \in A : S \notin \text{Sol}_f} \left|\, \tilde{S} \,\right| \leq \frac{2}{3} z.$$

We have

$$\text{OPT} = |\mathcal{C}(A)| = \left| \bigcup_{S \in A : S \notin \text{Sol}_f} \tilde{S} \right| + \left| \bigcup_{S \in A : S \in \text{Sol}_f} S \right| \leq \sum_{S \in A : S \notin \text{Sol}_f} \left|\tilde{S}\right| + |\mathcal{C}(\text{Sol}_f)| \leq \frac{2}{3} z + |\mathcal{C}(\text{Sol}_f)| .$$

Finally, the above implies $|\mathcal{C}(\text{Sol}_f)| \geq \text{OPT}/3$. □

Lemma 14. *Suppose step $(*)$ occurred to set S, then*

$$\max\{|S|\,, |\mathcal{C}(\text{Sol})|\} \geq (1 - \epsilon)\,\text{OPT}\,/3.$$

Proof. Using the analysis of Lemma 12, though S is not added to Sol, we still have

$$|S \setminus \mathcal{C}(\text{Sol})| + |\mathcal{C}(\text{Sol})| \geq \frac{2z}{3} \cdot \frac{w(\{S\}) + w(\text{Sol})}{L} \geq \frac{2z}{3}.$$

The seconds inequality holds due to the fact that $w(\{S\}) + w(\text{Sol}) > L$. Then, the claim immediately follows. □

Therefore, given a near correct guess $z \approx_{\epsilon^-} \text{OPT}$, the algorithm returns a $1/3 - \epsilon$ approximation. Similar to previous algorithms, we run this algorithm with different guesses $z = (1+\epsilon)^0, (1+\epsilon)^1, (1+\epsilon)^2, \ldots$. Finally, we return the solution with the maximum coverage. Each algorithm instantiation uses $\tilde{O}(n)$ space. Therefore, we have proved the following.

Theorem 15. *There exists a single-pass, $\tilde{O}(\epsilon^{-1}n)$-space algorithm that finds a $1/3 - \epsilon$ approximation of budgeted* MaxSetCoverage.

3 Algorithms for Maximum k-Vertex Coverage

In this section, we present algorithms for the maximum k-vertex coverage problem.

Notations. Given a graph G and a subset of nodes S, we define $\mathcal{C}_G(S)$ to be the number of edges with at least one endpoint in S. Recall that the maximum k-vertex coverage problem is to approximate the maximum value of $\mathcal{C}_G(S)$ over all sets S containing k nodes. We use E_G and V_G to denote the set of edges and nodes of the graph G respectively. If S is a subset of V_G, then $E_G(S)$ is the set of edges of the subgraph induced by S.

A core idea to our approach is to use *graph sparsification*:

Definition 16 (ϵ-sparsifier)**.** *Given a graph $G = (V, E)$, we say that a weighted subgraph $H = (V, E')$ is an ϵ-sparsifier for G if for all $S \subseteq V$, $\delta_G(S) \approx_\epsilon \delta_H(S)$ where $\delta_G(S)$ is the size of the cut $(S, V \setminus S)$ in G and $\delta_H(S)$ is the total weight of the cut $(S, V \setminus S)$ in H.*

It is known that any graph on N nodes has an ϵ-sparsifier with only $\tilde{O}(\epsilon^{-2}N)$ edges [41]. Furthermore, an ϵ-sparsifier can be constructed in the dynamic graph stream model using one pass and $\tilde{O}(\epsilon^{-2}N)$ space [24, 26].

In the next lemma we observe that, in addition to preserving the sizes of cuts of a graph, an ϵ-sparsifier also preserves the coverage of all subset of vertices up to a $1 + \epsilon$ factor. Note that for a weighted graph H, we define $\mathcal{C}_H(S)$ to be the total weight of edges with at least one endpoint in S.

Lemma 17. *Let H be an ϵ-sparsifier of G. For all $S \subseteq V$, $\mathcal{C}_G(S) \approx_\epsilon \mathcal{C}_H(S)$.*

Proof. First note that

$$\mathcal{C}_G(S) = 1/2 \cdot (\delta_G(S) + \sum_{v \in S} \delta_G(v))$$

and similarly $\mathcal{C}_H(S) = 1/2 \cdot (\delta_H(S) + \sum_{v \in S} \delta_H(v))$. Then the fact that $\mathcal{C}_G(S) \approx_\epsilon \mathcal{C}_H(S)$ follows since $\delta_G \approx_\epsilon \delta_H$. □

The following theorem follows immediately from the above lemma and the fact that it is possible to efficiently construct an ϵ-sparsifier in a single pass.

Theorem 18. *There exists a single-pass, $\tilde{O}(\epsilon^{-2}N)$-space algorithm that finds a $1 - \epsilon$ approximation of* MaxVertexCoverage *with high probability.*

Note that the above theorem assumes unbounded post-processing time. If k is constant, the post-processing will be polynomial. For larger k, if we still require polynomial running time then, after constructing the ϵ-sparsifier H, we could either use the 3/4 approximation algorithm via linear programming [1] or the folklore $(1 - 1/e)$ approximation greedy algorithm.

3.1 Algorithm for Near-Regular Graphs

In this section, we show that is possible to reduce the space used to $\tilde{O}(k)$ in the case of graphs that are regular or nearly regular. Define $\kappa \leq 1$ to be the ratio between the smallest degree and the largest degree; for a regular graph $\kappa = 1$. We show that a $(\kappa - \epsilon)$ approximation is possible using $\tilde{O}(\epsilon^{-3}k)$ space. This also implies a $(1 - \epsilon)$ approximation for regular graphs.

Theorem 19. *There exists a single-pass, $\tilde{O}(\epsilon^{-3}k)$-space algorithm that finds a $(\kappa - \epsilon)$ approximation of* MaxVertexCoverage *with high probability.*

Proof. **Case 1:** $N > 2k/\epsilon$. Let the minimum degree be t_1 and the maximum degree be $t_2 = t_1/\kappa$. We sample k nodes S uniformly random and note that the expected number of edges amongst these nodes is

$$\begin{aligned}
\mathrm{E}\left[|E_G(S)|\right] &= \frac{1}{2}\sum_{i=1}^{N} \sum_{j:(i,j)\in E_G} I_{i,j\in S} = \frac{1}{2}\sum_{i=1}^{N} \sum_{j:(i,j)\in E_G} \Pr\left[j \in S \mid i \in S\right] \Pr\left[i \in S\right] \\
&= \frac{1}{2}\sum_{i=1}^{N} \sum_{j:(i,j)\in E_G} \frac{(k-1)k}{(N-1)N} \leq \frac{1}{2}\frac{t_2 k^2}{N-1} < \frac{t_2 k^2}{N} \leq \frac{\epsilon t_2 k}{2}.
\end{aligned}$$

Therefore, by an application of Markov's inequality, $\Pr[|E_G(S)| \geq \epsilon k t_2] \leq 1/2$. Thus, if we sample $O(\log N)$ such set of k nodes in parallel, with high probability, there is a sample set S of k nodes satisfying $|E_G(S)| < kt_2\epsilon$ which implies

$$\mathcal{C}_G(S) \geq kt_1 - \epsilon k t_2 \geq k t_2 (t_1/t_2 - \epsilon) \geq (t_1/t_2 - \epsilon)\,\mathrm{OPT} = (\kappa - \epsilon)\,\mathrm{OPT}\,.$$

Case 2: $N \leq 2k/\epsilon$. We construct an ϵ-sparsifier algorithm for G. The required space is $\tilde{O}(\epsilon^{-3}k)$. As argued in Lemma 17, the coverage of all subset of nodes are preserved up to a $1+\epsilon$ factor. □

3.2 Extension to hypergraphs

We extend the above results to hypergraphs. The size of a cut $(S, V \setminus S)$ in a hypergraph G is defined as the number of hyperedges that contain at least one node in both S and $V \setminus S$. Again, we use $\delta_G(S)$ to denote the size of the cut. In the case that G is weighted, $\delta_G(S)$ denotes the total weight of the cut. We define an ϵ-sparsifier of a hypergraph similarly as in normal graphs. The sparsifier of a hypergraph in which the size of the hyperedges is bounded by d (rank d hypergraphs) can be constructed in the dynamic graph stream model using one pass and $\tilde{O}(\epsilon^{-2}dN)$ space [24].

Given a set of nodes S, we define the coverage of S as the number of hyperedges that have at least one node in S. First, we show that it is possible to approximate all the coverages by constructing a sparsifier of a slightly modified graph. In particular, we construct the sparsifier of the graph G' with an extra node v, i.e., $V_{G'} = V_G \cup \{v\}$, and for every hyperedge $e \in E_G$, we put the hyperedge $e \cup \{v\}$ in $E_{G'}$. It is easy to see that for all S that is a subset of V_G,

$$\mathcal{C}_G(S) = \delta_{G'}(S).$$

Therefore, it is immediate that we could $1+\epsilon$ approximate all the coverages in G by constructing the sparsifer of G'.

Theorem 20. *There exists a single-pass, $\tilde{O}(\epsilon^{-2}dN)$-space algorithm that finds a $1-\epsilon$ approximation of* MaxVertexCoverage *of rank d hypergraphs with high probability.*

Similarly, we develop an algorithm for hypergraphs with bounded κ^{-1}.

Theorem 21. *There exists a single-pass, $\tilde{O}(\epsilon^{-3}dk)$-space algorithm that finds a $(\kappa-\epsilon)$ approximation of* MaxVertexCoverage *of hypergraphs of rank d with high probability .*

Proof. Suppose we uniformly sample a set S of k nodes. If a hyperedge y is covered by S, then let $L_S(y)$ be the number of nodes in both y and S. The coverage of S satisfies

$$\mathcal{C}_G(S) \geq kt_1 - \sum_{y \in E_G(S)} (L_S(y) - 1) = kt_1 - \sum_{y \in E_G} L'_S(y).$$

where

$$L'_S(y) = \begin{cases} L_S(y) - 1 & \text{if } y \text{ is covered by } S \\ 0 & \text{otherwise} \end{cases}.$$

One can interpret the above inequality as follows. Each node in S covers at least t_1 hyperedges. Besides, we also need to subtract the number of times the covered hyperedges are over-counted which is $\sum_{y \in E_G(S)} (L_S(y) - 1)$. Let $\xi_y(j)$ denote the event that j nodes in the hyperedge y are in S and let $|y|$ denote the number of nodes in y. We have

$$\begin{aligned} \mathrm{E}[L'_S(y)] &= \sum_{j=1}^{|y|} (j-1)\Pr[\xi_y(j)] = \sum_{j=0}^{|y|} (j-1)\Pr[\xi_y(j)] + \Pr[\xi_y(0)] \\ &= \left(\sum_{j=0}^{|y|} j \Pr[\xi_y(j)] \right) - 1 + \Pr[\xi_y(0)]\,. \end{aligned}$$

The sum $\sum_{j=0}^{|y|} j \Pr[\xi_y(j)]$ is the expected value of the hypergeometric distribution and therefore it evaluates to $|y|k/N$. Furthermore,

$$\begin{aligned}\Pr[\xi_y(0)] &= \left(1-\frac{|y|}{N}\right)\left(1-\frac{|y|}{N-1}\right)\cdots\left(1-\frac{|y|}{N-k+1}\right)\\ &\leq \left(1-\frac{|y|}{N}\right)^k \leq \exp\left(-\frac{k|y|}{N}\right) \leq 1-\frac{k|y|}{N}+\frac{1}{2}\left(\frac{k|y|}{N}\right)^2.\end{aligned}$$

The last inequality follows from taking the first three terms of the Taylor's expansion. Putting it together,

$$\mathrm{E}[L'_S(y)] \leq \frac{k|y|}{N} - 1 + 1 - \frac{k|y|}{N} + \frac{1}{2}\left(\frac{k|y|}{N}\right)^2 = \frac{1}{2}\left(\frac{k|y|}{N}\right)^2.$$

Hence, if $N \geq 4kd/\epsilon$, then

$$\sum_{y\in E_G} \mathrm{E}[L'_S(y)] \leq \frac{1}{2}\sum_{y\in E_G}\left(\frac{k|y|}{N}\right)^2 \leq \frac{1}{2}d\left(\frac{k}{N}\right)^2 \sum_{y\in E_G}|y| \leq \frac{1}{2}d\left(\frac{k}{N}\right)^2 Nt_2 \leq \frac{1}{8}\epsilon k t_2.$$

By an application of Markov's inequality, $\Pr\left[\sum_{y\in E_G(S)} L'_S(y) \geq \epsilon k t_2\right] \leq 1/8$. Thus, if we sample $O(\log N)$ such set of k nodes in parallel, with high probability, there is a sample set S of k nodes satisfying $\sum_{y\in E_G(S)} L'_S(y) \leq \epsilon k t_2$ which implies that $\mathcal{C}_G(S) \geq kt_1 - \epsilon k t_2 \geq (\kappa-\epsilon)\,\mathrm{OPT}$. If $N \leq 4kd/\epsilon$, we simply construct the sparsifier of G' as described above to achieve a $1-\epsilon$ approximation. □

4 Lower Bounds

In this section, we prove space lower bounds for data stream algorithms that approximate MaxSetCoverage or MaxVertexCoverage. In particular, these imply that improving over an $(1-1/e)$ approximation of MaxSetCoverage with constant passes and constant k, requires $\Omega(m)$ space. Recall that, still assuming k is constant, we designed a constant-pass algorithm that returned a $(1-1/e-\epsilon)$ approximation using $\tilde{O}(\epsilon^{-2}k)$ space. For constant k, we also show that improving over a κ approximation (where κ is the ratio between the lowest degree and the highest degree) for MaxVertexCoverage requires $\Omega(N\kappa^3)$ space. Our algorithm returned a $\kappa-\epsilon$ approximation using $\tilde{O}(\epsilon^{-3}k)$ space.

Approach. We prove both bounds by a reduction from r-player set-disjointness in communication complexity. In this problem, there are r players where the ith player has a set $S_i \subseteq [u]$. It is promised that exactly one of the following two cases happens.

- Case 1 (NO instance): All the sets are pairwise disjoint.
- Case 2 (YES instance): There is a unique element $e \in [u]$ such that $e \in S_i$ for all $i \in [r]$.

The goal of the communication problem is the rth player answers whether the input is a YES instance or a NO instance correctly with probability at least 0.9. We shall denote this problem by $\mathsf{DISJ}_r(u)$.

The communication complexity of the above problem in p-round, one-way model (where each round consists of player 1 sending a message to player 2, then player 2 sending a message to player 3 and so on) is $\Omega(u/r)$ [16] even if the players may use public randomness. This implies that in any randomized communication protocol, the maximum message sent by a player contains $\Omega(u/(pr^2))$ bits. Without loss of generality, we could assume that $|S_1 \cup S_2 \cup \ldots \cup S_r| \geq u/4$ via a padding argument.

Theorem 22. *Assuming $n = \Omega(\epsilon^{-2}k\log m)$, any constant-pass algorithm that finds a $(1+\epsilon)(1-(1-1/k)^k)$ approximation of* MaxSetCoverage *with probability at least 0.99 requires $\Omega(m/k^2)$ space even when all the sets have the same size.*

Recall that $(1+\epsilon)(1-(1-1/k)^k) \geq 1-1/e+O(\epsilon)$.

Proof. Our proof is a reduction from $\mathsf{DISJ}_k(m)$. Consider a sufficiently large n where k divides n. For each $i \in [m]$, let $\mathcal{P}_i$ be a random partition of $[n]$ into k sets $V_1^i, \ldots, V_k^i$ of equal size. Each partition is chosen independently and the players agree on these partitions using public randomness before receiving the input.

For each player j, if $i \in S_j$, then she puts V_j^i in the stream. According to the aforementioned assumption, the stream consists of at least $m/4$ sets.

If the input is a NO instance, then for each $i \in [m]$, there is at most one set V_j^i in the stream. Hence, the stream consists of independent random sets of size n/k. Therefore, for each $e \in [n]$ and any k sets $V_{j_1}^{i_1}, \ldots, V_{j_k}^{i_k}$ in the stream,

$$\Pr\left[e \in V_{j_1}^{i_1} \cup \ldots \cup V_{j_k}^{i_k}\right] = 1 - (1 - 1/k)^k.$$

By an application of Chernoff bound for negatively correlated boolean random variables [37],

$$\Pr\left[\left||V_{j_1}^{i_1} \cup \ldots \cup V_{j_k}^{i_k}| - \left(1 - \left(1 - \frac{1}{k}\right)^k\right) n\right| > \epsilon \left(1 - \left(1 - \frac{1}{k}\right)^k\right) n\right]$$
$$\leq 3\exp\left(-\epsilon^2 \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \frac{n}{3}\right) \leq 3\exp\left(-\epsilon^2(1 - 1/e)n/3\right) \leq \frac{1}{m^{10+k}}.$$

The last inequality holds when n is a sufficiently large multiple of $k\epsilon^{-2}\log m$. Therefore, the maximum coverage in this case is at most $(1+\epsilon)(1-(1-1/k)^k)n$ with probability at least $1 - 1/m^{10}$ by taking the union bound over all $\binom{m}{k} \leq m^k$ possible k sets.

If the input is a YES instance, then clearly, the maximum coverage is n. This is because there exists $i \in [m]$ such that $i \in S_1 \cap \ldots \cap S_k$ and therefore $V_1^i, \ldots, V_k^i$ are in the stream.

Therefore, any constant pass and $O(s)$-space algorithm that finds a $(1+2\epsilon)(1-(1-1/k)^k)$ approximation of the maximum coverage with probability at least 0.99 implies a protocol to solve the k-party disjointness problem using $O(s)$ bits of communication. Thus, $s = \Omega(m/k^2)$ as required. □

Consider the sets $S_1, \ldots, S_r \subseteq [u]$ that satisfy the unique intersection promise as in $\mathsf{DISJ}_r(u)$. Let X be the r by u matrix in which the row X_i is the characteristic vector of S_i. Suppose there are $r' = \Omega(r^2)$ players. Chakrabarti et al. [15] showed that if each entry of X is given to a unique player and the order in which the entries are given to the players is random, then the players need to use $\Omega(u/r)$ bits of communication to tell whether the sets is a YES instance or a NO instance with probability at least 0.9. Thus, in any randomized protocol, the maximum message sent by a player contains $\Omega(u/r^3)$ bits. Hence, using the same reduction and assuming constant k, we show that the same lower bound holds even for random order stream.

Theorem 23. *Assuming $n = \Omega(\epsilon^{-2}k\log m)$, any constant-pass algorithm that finds a $(1+\epsilon)(1-(1-1/k)^k)$ approximation of* MaxSetCoverage *with probability at least 0.99 requires $\Omega(m/k^3)$ space even when all the sets have the same size and arrive in random order.*

Next, we prove a lower bound for the k-vertex coverage problem for graphs where the ratio between the minimum degree and the maximum degree is at least κ. We show that for constant k, beating κ approximation for constant κ requires $\Omega(N)$ space. Since κ can be smaller than any constant, this also establishes that $\Omega(N)$ space is required for any constant approximation of MaxVertexCoverage.

Theorem 24. *For $\epsilon > 0$, any constant-pass algorithm that finds a $(\kappa+\epsilon)$ approximation of* MaxVertexCoverage *with probability at least 0.99 requires $\Omega(N\kappa^3/k)$ space.*

Proof. Initially, assume $k = 1$. We consider the multi-party set disjointness problem $\mathsf{DISJ}_t(N')$ where $t = 1/\kappa$ and $N' = N/t$. Here, there are t players and the input sets are subsets of $[N']$. We consider a bipartite graph where the set of possible nodes are $L \cup R$ where $L = \{u_i\}_{i\in[N']}$ and $R = \{v_{i,j}\}_{i\in[N'],j\in[t]}$. Note that this graph has $(t+1)N' = \Theta(N)$ nodes. However we only consider a node to exist if the stream contains an edge incident to that node.

The j-th player defines a set of edges on this graph based on their set S_j as follows. If $i \in S_j$ she puts the edge between u_i and $v_{i,j}$. If $S_1, \ldots, S_t$ is a YES instance, then there must be a node u_i that has degree t. If A is a NO instance, then every node in the graph has degree at most 1. Hence the ratio of minimum degree to maximum degree is at least $1/t = \kappa$ as required.

Thus, for $k = 1$, a $1/t$ approximation with probability at least 0.99 on a graph of N nodes implies a protocol to solve $\mathsf{DISJ}_t(N')$. Therefore, the algorithm requires $\Omega(N\kappa^3)$ space. For general k, we make k copies of the above construction to deduce the lower bound $\Omega(N\kappa^3/k)$.

□

References

[1] Alexander A. Ageev and Maxim Sviridenko. Approximation algorithms for maximum coverage and max cut with given sizes of parts. In *IPCO*, volume 1610 of *Lecture Notes in Computer Science*, pages 17–30. Springer, 1999.

[2] Kook Jin Ahn, Graham Cormode, Sudipto Guha, Andrew McGregor, and Anthony Wirth. Correlation clustering in data streams. In *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2237–2246. JMLR.org, 2015.

[3] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. In *SODA*, pages 459–467. SIAM, 2012.

[4] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketches: sparsification, spanners, and subgraphs. In *PODS*, pages 5–14. ACM, 2012.

[5] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Spectral sparsification in dynamic graph streams. In *APPROX-RANDOM*, volume 8096 of *Lecture Notes in Computer Science*, pages 1–10. Springer, 2013.

[6] Aris Anagnostopoulos, Luca Becchetti, Ilaria Bordino, Stefano Leonardi, Ida Mele, and Piotr Sankowski. Stochastic query covering for fast approximate document retrieval. *ACM Trans. Inf. Syst.*, 33(3):11:1–11:35, 2015.

[7] Sepehr Assadi, Sanjeev Khanna, and Yang Li. Tight bounds for single-pass streaming complexity of the set cover problem. In *STOC*, pages 698–711. ACM, 2016.

[8] Sepehr Assadi, Sanjeev Khanna, Yang Li, and Grigory Yaroslavtsev. Maximum matchings in dynamic graph streams and the simultaneous communication model. In *SODA*, pages 1345–1364. SIAM, 2016.

[9] Giorgio Ausiello, Nicolas Boria, Aristotelis Giannakos, Giorgio Lucarelli, and Vangelis Th. Paschos. Online maximum k-coverage. *Discrete Applied Mathematics*, 160(13-14):1901–1913, 2012.

[10] Ashwinkumar Badanidiyuru, Baharan Mirzasoleiman, Amin Karbasi, and Andreas Krause. Streaming submodular maximization: massive data summarization on the fly. In *KDD*, pages 671–680. ACM, 2014.

[11] MohammadHossein Bateni, Hossein Esfandiari, and Vahab S. Mirrokni. Almost optimal streaming algorithms for coverage problems. *CoRR*, abs/1610.08096, 2016.

[12] Sayan Bhattacharya, Monika Henzinger, Danupon Nanongkai, and Charalampos E. Tsourakakis. Space- and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams. In *STOC*, pages 173–182. ACM, 2015.

[13] Édouard Bonnet, Bruno Escoffier, Vangelis Th. Paschos, and Georgios Stamoulis. A 0.821-ratio purely combinatorial algorithm for maximum k-vertex cover in bipartite graphs. In *LATIN*, volume 9644 of *Lecture Notes in Computer Science*, pages 235–248. Springer, 2016.

[14] Bugra Caskurlu, Vahan Mkrtchyan, Ojas Parekh, and K. Subramani. On partial vertex cover and budgeted maximum coverage problems in bipartite graphs. In *IFIP TCS*, volume 8705 of *Lecture Notes in Computer Science*, pages 13–26. Springer, 2014.

[15] Amit Chakrabarti, Graham Cormode, and Andrew McGregor. Robust lower bounds for communication and stream computation. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:62, 2011.

[16] Amit Chakrabarti, Subhash Khot, and Xiaodong Sun. Near-optimal lower bounds on the multi-party communication complexity of set disjointness. In *IEEE Conference on Computational Complexity*, pages 107–117. IEEE Computer Society, 2003.

[17] Amit Chakrabarti and Anthony Wirth. Incidence geometries and the pass complexity of semi-streaming set cover. In *SODA*, pages 1365–1373. SIAM, 2016.

[18] Chandra Chekuri, Shalmoli Gupta, and Kent Quanrud. Streaming algorithms for submodular function maximization. In *ICALP (1)*, volume 9134 of *Lecture Notes in Computer Science*, pages 318–330. Springer, 2015.

[19] Rajesh Chitnis, Graham Cormode, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Andrew McGregor, Morteza Monemizadeh, and Sofya Vorotnikova. Kernelization via sampling with applications to finding matchings and related problems in dynamic graph streams. In *SODA*, pages 1326–1344. SIAM, 2016.

[20] Graham Cormode, Mayur Datar, Piotr Indyk, and S. Muthukrishnan. Comparing data streams using hamming norms (how to zero in). *IEEE Trans. Knowl. Data Eng.*, 15(3):529–540, 2003.

[21] Graham Cormode, Howard J. Karloff, and Anthony Wirth. Set cover algorithms for very large datasets. In *CIKM*, pages 479–488. ACM, 2010.

[22] Yuval Emek and Adi Rosén. Semi-streaming set cover - (extended abstract). In *ICALP (1)*, volume 8572 of *Lecture Notes in Computer Science*, pages 453–464. Springer, 2014.

[23] Uriel Feige. A threshold of ln n for approximating set cover. *J. ACM*, 45(4):634–652, 1998.

[24] Sudipto Guha, Andrew McGregor, and David Tench. Vertex and hyperedge connectivity in dynamic graph streams. In *PODS*, pages 241–247. ACM, 2015.

[25] Sariel Har-Peled, Piotr Indyk, Sepideh Mahabadi, and Ali Vakilian. Towards tight bounds for the streaming set cover problem. In *PODS*, pages 371–383. ACM, 2016.

[26] Michael Kapralov, Yin Tat Lee, Cameron Musco, Christopher Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. In *FOCS*, pages 561–570. IEEE Computer Society, 2014.

[27] Michael Kapralov and David P. Woodruff. Spanners and sparsifiers in dynamic streams. In *PODC*, pages 272–281. ACM, 2014.

[28] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. *Theory of Computing*, 11:105–147, 2015.

[29] Samir Khuller, Anna Moss, and Joseph Naor. The budgeted maximum coverage problem. *Inf. Process. Lett.*, 70(1):39–45, 1999.

[30] Dmitry Kogan and Robert Krauthgamer. Sketching cuts in graphs and hypergraphs. In *6th Innovations in Theoretical Computer Science*, 2015.

[31] Christian Konrad. Maximum matching in turnstile streams. In *ESA*, volume 9294 of *Lecture Notes in Computer Science*, pages 840–852. Springer, 2015.

[32] Andreas Krause and Carlos Guestrin. Near-optimal observation selection using submodular functions. In *AAAI*, pages 1650–1654. AAAI Press, 2007.

[33] Ravi Kumar, Benjamin Moseley, Sergei Vassilvitskii, and Andrea Vattani. Fast greedy algorithms in mapreduce and streaming. *TOPC*, 2(3):14, 2015.

[34] Andrew McGregor. Graph stream algorithms: a survey. *SIGMOD Record*, 43(1):9–20, 2014.

[35] Andrew McGregor, David Tench, Sofya Vorotnikova, and Hoa T. Vu. Densest subgraph in dynamic graph streams. In *MFCS (2)*, volume 9235 of *Lecture Notes in Computer Science*, pages 472–482. Springer, 2015.

[36] Andrew McGregor, Sofya Vorotnikova, and Hoa T. Vu. Better algorithms for counting triangles in data streams. In *PODS*, pages 401–411. ACM, 2016.

[37] Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff-hoeffding bounds. *SIAM J. Comput.*, 26(2):350–368, 1997.

[38] Jaikumar Radhakrishnan and Saswata Shannigrahi. Streaming algorithms for 2-coloring uniform hypergraphs. In *Algorithms and Data Structures - 12th International Symposium, WADS 2011, New York, NY, USA, August 15-17, 2011. Proceedings*, pages 667–678, 2011.

[39] Barna Saha and Lise Getoor. On maximum coverage in the streaming model & application to multi-topic blog-watch. In *SDM*, pages 697–708. SIAM, 2009.

[40] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM J. Discrete Math.*, 8(2):223–250, 1995.

[41] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, 2011.

[42] He Sun. Counting hypergraphs in data streams. *CoRR*, abs/1304.7456, 2013.

[43] Huiwen Yu and Dayu Yuan. Set coverage problems in a one-pass data stream. In *SDM*, pages 758–766. SIAM, 2013.

A $1/2 - \epsilon$ approximation in one pass

We rephrase the algorithm of Badanidiyuru et al. [10] in the context of maximum coverage in Algorithm 4. The additional detail is to maintain the F_0 sketch $\mathcal{T}$ of Sol. Given a near-correct guess $z \approx_{\epsilon^-} \mathrm{OPT}$, this single-pass algorithm uses $\tilde{O}(z+k)$ space and finds a $1/2-\epsilon$ approximation.

Let c be some sufficiently large constant. We could assume that $k \le \epsilon^2 n/(8c\log m)$. Otherwise, we already have an algorithm that uses $O(n\log n + k\log m) = \tilde{O}(\epsilon^{-2}k)$ space assuming $z \approx_{\epsilon^-} \mathrm{OPT}$. Hence, by trying different guesses $z = (1+\epsilon)^0, (1+\epsilon)^1, (1+\epsilon)^2, \ldots$, we achieve a single-pass, $\tilde{O}(\epsilon^{-3}k)$-space algorithm that finds a $1/2 - O(\epsilon)$ approximation as required.

Since $k > \epsilon^2 n/(8c\log m)$, we could reuse the subsampling strategy with the guesses $v = (1+\epsilon)^0, (1+\epsilon)^1, (1+\epsilon)^2, \ldots$ to yield Theorem 11. In particular, when $v \approx_{2\epsilon} \mathrm{OPT}$, we have

$$\mathrm{OPT}' = (1\pm\epsilon)p\,\mathrm{OPT} = (1\pm\epsilon)\frac{c\epsilon^{-2}k\log m}{(1\pm 2\epsilon)\,\mathrm{OPT}}\,\mathrm{OPT} \approx_{4\epsilon} c\epsilon^{-2}k\log m$$

according to Corollary 6. Therefore, by setting $z = (1-\epsilon)(1-4\epsilon)c\epsilon^{-2}k\log m$, we know that $z \approx_{(5\epsilon)^-} \mathrm{OPT}'$ when $v \approx_{2\epsilon} \mathrm{OPT}$.

Applying Algorithm 4 and Corollary 6 to the subsampled elements corresponding to the guess $v \approx_{2\epsilon} \mathrm{OPT}$, we have an algorithm instantiation that uses $\tilde{O}(\epsilon^{-2}k)$ space and finds a $1 - 1/e - O(\epsilon)$ approximation.

The algorithm is provided $z \approx_{\epsilon^-} \mathrm{OPT}$.
$C \leftarrow \emptyset$, $\mathrm{Sol} \leftarrow \emptyset$, $\mathcal{T} \leftarrow \mathcal{M}(\emptyset)$. // $\mathcal{T}$ maintains the F_0 sketch of Sol
During the stream, on seeing a set S:
  Store the F_0 sketch $\mathcal{M}(S)$.
  if $|S \setminus C| \geq (z/2 - |C|)/(k - |\,\mathrm{Sol}\,|)$ **then**
    if $|\,\mathrm{Sol}\,| < k$ **then**
      $\mathrm{Sol} \leftarrow \mathrm{Sol} \cup \{S\}$. // Sol is only used for the analysis.
      $\mathrm{ID}_{\mathrm{Sol}} \leftarrow \mathrm{ID}_{\mathrm{Sol}} \cup \{\mathrm{ID}(S)\}$.
      $C \leftarrow C \cup S$.
      $\mathcal{T} \leftarrow \mathcal{T} + \mathcal{M}(\mathrm{Sol} \cup S)$. // update the F_0 sketch of Sol
    end
    if $|\,\mathrm{Sol}\,| = k$ **then** stop the algorithm and **return** $\mathrm{ID}_{\mathrm{Sol}}$.
  end
Post-processing: **return** $\mathrm{ID}_{\mathrm{Sol}}$.

Algorithm 4: The basic algorithm for a $1/2 - \epsilon$ approximation in one pass and $\tilde{O}(\epsilon^{-3}k)$ space.